Measurement errors in network data

M. E. J. Newman
Department of Physics and Center for the Study of Complex Systems, University of Michigan, Ann Arbor, Michigan, USA

The recent growth in interest in the physics and mathematics of networks has been driven in large part by the increasing availability of data describing the structure of networks ranging from the internet and the web to social and biological networks. It is a surprising feature of many empirical studies, however, that the data are reported without any estimate of their expected accuracy, even though it is clear that most do suffer from measurement error of various kinds. In this paper we develop the theory of measurement error for network data and give an expectation–maximization algorithm for estimating both false positive and false negative rates for edges in observed networks. We give an example application of our methods to social networks determined from proximity data. The methods we describe are general, and could be extended straightforwardly to cover different types of networks or errors.

arXiv:1703.07376v1 [cs.SI] 21 Mar 2017

There has been substantial interest in the last two decades in the form and function of networked systems, including technological, biological, and social networks. This interest has been bolstered by the availability of many new data sets, some of them very extensive, describing the structure of networks in different areas of human endeavor and the natural world. A striking, and potentially problematic, feature of many of these data sets, however, is that they include little or no information on the reliability of the data they contain. In essentially every other field of science it is required practice when publishing the results of empirical studies to give error estimates, to allow others to gauge the reliability of measurements. Network data, by contrast, are rarely published with error estimates, and most analyses simply assume that the measured structure of a network is the true structure, free of experimental error.

Yet in virtually no case is this actually true. Almost all studies of real-world network structures contain experimental error of some kind, and frequently of many kinds simultaneously. Studies of social networks suffer from subjectivity on the part of participants (and potentially also experimenters), quantification, coding, and recording error, and missing data [1–3]. Biological networks suffer from laboratory experimental error as well as recording errors and publication bias [4, 5]. Technological networks suffer from sampling bias, technical limitations, and missing data [6, 7]. Moreover, simulation and other studies suggest that errors in network data can have substantial deleterious effects on our estimates of network metrics [8–12].

The issue of errors in network data has particularly been recognized in the social sciences, where there has been extensive discussion of sources of error in social surveys, its effects on measurements, and ways of minimizing it [1, 3, 13–15]. There is also substantial domain-specific literature on problems such as predicting missing nodes or edges in networks [16–20] and name disambiguation in bibliometrics [21–24], both of which fall under the broader umbrella of network reconstruction methods [18, 25]. Here we take a complementary approach, focusing on the problem of parametrizing and estimating errors, giving both a general theory of network error as well as specific methods for the case of independent edge errors. In the process, we also aim to highlight the importance of measurement error in network data, in the hope of encouraging greater attention to these issues in future.

For the sake of a focused discussion, we will concentrate here on unweighted undirected networks. Errors in such networks can be of several kinds [3, 11] but divide broadly into node errors and edge errors. Possible errors in the nodes of a network include missing nodes (which are common in many networks) and extraneous nodes (nodes that are present in the data but don't really exist). The latter are rarer, except for one common special case, that of erroneous node aggregation and disaggregation. In social networks, for instance, a single person may be erroneously represented by two or more different nodes, perhaps because they go by different names in different contexts. Similarly two distinct people may be erroneously conflated and represented by a single node because they happen to have the same name.

Errors on edges, which will be our primary concern in this paper, include omitted edges (connections that exist in reality but are not recorded in the data), extraneous edges (connections that don't exist but *are* recorded in the data), and missing data (edges for which we have no information about their existence because, for instance, they were never measured). Missing data are often treated as if they were non-edges, but the two are not the same: absence of evidence is not evidence of absence.

Before developing the theory of these errors, let us briefly revisit the theory of statistical error in ordinary empirical measurements, such as measurements of a continuous real-valued numerical quantity. Suppose an experimenter measures a quantity of interest, for instance in the lab. Normally they will measure it more than once, which is necessary to get an estimate of error. The result

is a series of measurements $x_1, \ldots, x_N$, usually assumed to be independent. It is further assumed that there is a single true underlying value z of the quantity, sometimes called the *ground truth*, and that the measured data x_i are related to this value in some way that is captured by an *error model*. The error model tells us the probability $P(x|z,\theta)$ that we make a measurement x, given the ground truth z and (optionally) some additional model parameters, here denoted θ. By far the most common error model for real-valued data is a normal or Gaussian distribution, which has one parameter in addition to the ground truth, the standard deviation σ of the distribution.

Our goal, given the observed data, is to estimate z (the thing we are trying to measure) plus the parameters of the model, which, in combination with the model itself, quantify the error on our measurement. Typically estimation is done by maximum-likelihood methods. Applying such methods to the Gaussian model, for instance, leads to the standard formulas $z = (1/N)\sum_i x_i$ for the expected value of the ground truth and $\sigma^2 = (1/N)\sum_i (x_i - z)^2$ for the standard deviation. When one publishes an experimental result in the form of a measured value plus or minus an error, one is effectively declaring that one has assumed a Gaussian model for the error and giving the parameters of that model.

Our approach to quantifying experimental errors on network data is analogous. Given some measurements of the structure of a network we will assume an error model for the relationship between those measurements and the underlying ground truth, then use maximum-likelihood methods to infer both the ground truth and the parameters of the model from the data. We focus here specifically on errors on edges, which are arguably the most common type of network error, although the methods described could be extended to node errors as well.

Suppose then that we are interested in the structure of a certain n-node network and let us denote the true, ground-truth structure of that network by an $n \times n$ symmetric adjacency matrix $\mathbf{A}$, having elements $A_{ij} = 1$ if nodes i and j are connected by an edge and 0 otherwise. We make a set of measurements of the structure of this network, measurements that could potentially take many forms. In some cases, such as social networks, we might measure the presence or absence of edges between individual node pairs; in others, such as biological networks or the internet, we might measure complete paths through the network; while in others still, such as so-called affiliation networks, the unit of measurement might be a clique or subgraph within the larger network. In the process of making our measurements we may observe some edges many times, others not at all, but any of our measurements could potentially be incorrect because of experimental error. Moreover, the errors need not be independent: it is entirely possible for them to be correlated, potentially in complex ways.

Whatever form our data take, let us collectively denote them by D. Then the error model is a probability distribution $P(D|\mathbf{A},\theta)$ that defines the probability of making measurements D given the true structure of the network $\mathbf{A}$ plus model parameters θ. Our goal is estimate both $\mathbf{A}$ and θ, the latter characterizing the error on our measurements in a manner analogous to the standard deviation in the Gaussian model. We will estimate the parameters first, followed by the network structure.

Applying Bayes rule, we write

$$P(\mathbf{A},\theta|D) = \frac{P(D|\mathbf{A},\theta)P(\mathbf{A})P(\theta)}{P(D)}, \tag{1}$$

where it is axiomatic that the prior probability $P(\mathbf{A})$ on the network structure is independent of the error model and its parameters θ. Summing over all possible network structures $\mathbf{A}$, we get $P(\theta|D) = \sum_{\mathbf{A}} P(\mathbf{A},\theta|D)$, which we then maximize to find the most probable value of the parameters θ given the observed data. In fact, for convenience, we maximize not $P(\theta|D)$ but its logarithm, whose maximum falls in the same place.

Employing the well-known Jensen inequality $\log \sum_i x_i \geq \sum_i q_i \log(x_i/q_i)$, we can write

$$\log P(\theta|D) = \log \sum_{\mathbf{A}} P(\mathbf{A},\theta|D) \geq \sum_{\mathbf{A}} q(\mathbf{A}) \log \frac{P(\mathbf{A},\theta|D)}{q(\mathbf{A})}, \tag{2}$$

where $q(\mathbf{A})$ is any probability distribution over networks $\mathbf{A}$ satisfying $\sum_{\mathbf{A}} q(\mathbf{A}) = 1$. It is trivially the case that exact equality between left- and right-hand sides of Eq. (2) is achieved when

$$q(\mathbf{A}) = \frac{P(\mathbf{A},\theta|D)}{\sum_{\mathbf{A}} P(\mathbf{A},\theta|D)}, \tag{3}$$

and hence this choice maximizes the right-hand side with respect to q. A further maximization with respect to θ will then give us the maximum-likelihood value we seek. To put that another way, a double maximization of the right-hand side of (2) with respect to both q and θ will give us our answer for θ.

At first glance this does not seem a promising approach: it transforms what was a simple maximization of $P(\theta|D)$ with respect to θ alone into a double maximization over both θ and q. In fact, however, it turns out to be highly convenient. The double maximization can be easily carried out by maximizing first with respect to $q(\mathbf{A})$ using Eq. (3) and then with respect to θ, repeating until the result converges. The maximization with respect to θ can be achieved by simple differentiation. Differentiating (2) while holding $q(\mathbf{A})$ constant, then setting the result to zero, we derive the condition

$$\sum_{\mathbf{A}} q(\mathbf{A}) \nabla_\theta \log P(\mathbf{A},\theta|D) = 0. \tag{4}$$

The solution of this equation gives us our value for θ.

Our method consists of iterating Eqs. (3) and (4) from random initial values to convergence. The final result is a value for the parameters θ, which, in combination with the error model itself, gives us our error estimate. In the process, we also get an estimate of the ground-truth network structure. Indeed, we get the entire posterior probability distribution over structures, since from Eq. (3) the quantity $q(\mathbf{A})$ is none other than $q(\mathbf{A}) = P(\mathbf{A},\theta|D)/P(\theta|D) = P(\mathbf{A}|D,\theta)$. In other words it is precisely the probability of the network having true structure $\mathbf{A}$ given the observed data and our estimate of the parameters θ. (This approach differs from other methods for generating distributions over networks, including our own previous work, that do not make use of an error model but instead make stronger assumptions about the particular type of structure the network contains [17, 18].)

The calculation given here is an example of an expectation–maximization or EM algorithm [26, 27]. The EM algorithm has been previously applied to a number of other problems in the theory of networks, including link prediction [20] and community detection [28]. Let us see how it is applied in practice to a specific error model.

Perhaps the simplest model we can adopt for edge error in networks is one in which we make measurements of the presence or absence of edges between individual node pairs, and those measurements are assumed to be independent identically distributed random variables, conditioned on the ground truth A_{ij} for the node pair i,j in question. That is, the probability of measuring an edge between nodes i and j depends only on A_{ij} and in the same way for all i,j. This dependence can be parametrized by two quantities: the *true positive rate* α which is the probability of observing an edge where one truly exists, and the *false positive rate* β, the probability of observing an edge where none exists. (The resulting model is similar to models for certain non-network problems involving binary misclassification error [29], such as assessing the accuracy of responses to test questions when the true answers are unknown [30, 31], although our methods of analysis are different.)

To estimate error rates it is necessary (just as with ordinary real-valued data) that we measure at least some of the edges in the network more than once. Suppose that for node pair i,j we make N_{ij} measurements and observe an edge to be present in E_{ij} of those measurements. Then under the independent edge model

$$P(D|\mathbf{A},\theta) = \prod_{i<j} \bigl[\alpha^{E_{ij}}(1-\alpha)^{N_{ij}-E_{ij}}\bigr]^{A_{ij}} \times \prod_{i<j} \bigl[\beta^{E_{ij}}(1-\beta)^{N_{ij}-E_{ij}}\bigr]^{1-A_{ij}}. \quad (5)$$

We will assume that the prior probability ρ of existence of all edges is the same, so that $P(\mathbf{A}|\rho) = \prod_{i<j} \rho^{A_{ij}}(1-\rho)^{1-A_{ij}}$ (although a different assumption would be straightforward to incorporate), and that the prior probability distributions on α, β, and ρ are all uniform in the interval $[0,1]$. Combining Eqs. (1) and (5) and taking logs, we then have

$$\begin{aligned}\log P(\mathbf{A},\theta|D) = {} & \\ & \textstyle\sum_{i<j} \bigl[E_{ij}A_{ij}\log\alpha + (N_{ij}-E_{ij})A_{ij}\log(1-\alpha) \\ & \quad + E_{ij}(1-A_{ij})\log\beta + (N_{ij}-E_{ij})(1-A_{ij})\log(1-\beta) \\ & \quad + A_{ij}\log\rho + (1-A_{ij})\log(1-\rho)\bigr] - \log P(D). \end{aligned} \quad (6)$$

Substituting into Eq. (4) and performing the derivatives, we arrive at expressions for the three model parameters α, β, and ρ:

$$\alpha = \frac{\sum_{i<j} E_{ij}Q_{ij}}{\sum_{i<j} N_{ij}Q_{ij}}, \quad \beta = \frac{\sum_{i<j} E_{ij}(1-Q_{ij})}{\sum_{i<j} N_{ij}(1-Q_{ij})}, \quad (7)$$

$$\rho = \frac{1}{\binom{n}{2}} \sum_{i<j} Q_{ij}, \quad (8)$$

where $Q_{ij} = \sum_{\mathbf{A}} A_{ij} q(\mathbf{A}) = P(A_{ij}=1|D,\theta)$ is the (posterior) probability that there is an edge between i and j. Combining Eqs. (1), (3), and (5), we find that

$$Q_{ij} = \frac{\rho\alpha^{E_{ij}}(1-\alpha)^{N_{ij}-E_{ij}}}{\rho\alpha^{E_{ij}}(1-\alpha)^{N_{ij}-E_{ij}} + (1-\rho)\beta^{E_{ij}}(1-\beta)^{N_{ij}-E_{ij}}}. \quad (9)$$

Note that if we make no measurements for a pair of nodes i,j so that $N_{ij} = E_{ij} = 0$ (the case of "missing data"), this expression correctly gives Q_{ij} equal to the prior edge probability ρ.

Our complete calculation now involves iterating Eqs. (7) and (8) and Eq. (9) until convergence is reached. In the calculations presented below, we start the iteration from random initial values of α uniform in $[\frac{1}{2},1]$, β uniform in $[0,\frac{1}{2}]$, and ρ equal to the empirical edge probability $m/\binom{n}{2}$, where m is the number of node pairs with $E_{ij} > 0$.

Armed with the results of these calculations, one would conventionally then report the most likely structure for the ground-truth network (which is given by placing an edge between every pair of nodes for which $Q_{ij} > \frac{1}{2}$) plus the parameters α and β. This would be the logical equivalent of the expected value and standard error we give for ordinary scalar measurements. The values of α and β parametrize the error on the reported structure: β tells us the probability of a false positive edge and $1-\alpha$ tells us the probability of a false negative.

In the present case, however, a more informative way to report the error would be to give the complete matrix of posterior edge probabilities Q_{ij}, which would allow others to reconstruct the entire posterior probability distribution over possible ground-truth networks:

$$P(\mathbf{A}|D,\theta) = \prod_{i>j} Q_{ij}^{A_{ij}}(1-Q_{ij})^{1-A_{ij}}. \quad (10)$$

The posterior distribution (for this or any error model) lets us compute the distribution of any other network quantity we might be interested in—degrees, correlations, clustering coefficients, and so forth. For any quantity $X(\mathbf{A})$ that is a function of the adjacency matrix $\mathbf{A}$, the probability distribution is

$$P(X = x|D,\theta) = \sum_{\mathbf{A}} \delta_{x,X(\mathbf{A})} P(\mathbf{A}|D,\theta), \qquad (11)$$

where δ_{ij} is the Kronecker delta. Alternatively, we could calculate the mean μ_X and variance σ_X^2 of X from

$$\mu_X = \sum_{\mathbf{A}} X(\mathbf{A}) P(\mathbf{A}|D,\theta), \qquad (12)$$

$$\sigma_X^2 = \sum_{\mathbf{A}} [X(\mathbf{A}) - \mu_X]^2 P(\mathbf{A}|D,\theta), \qquad (13)$$

although this is appropriate only when X has a normal or approximately normal distribution. The values of the Q_{ij} can also be used as inputs to further inference calculations, for instance of community structure [32].

One can think of the matrix $\mathbf{Q}$ of probabilities as a generalization of the conventional adjacency matrix. If $Q_{ij} = 0$ or 1 it behaves like a standard adjacency matrix element: there definitely is, or is not, an edge between i and j. For values in between it interpolates between these extremes. There are a few instances in the literature, particularly for biological networks, where experimenters have given error estimates on network structure in the form of posterior probabilities of this sort on edges [4, 5], although estimated using different, domain-specific methods. Such examples are, however, rare. Indeed, there are relatively few studies that have even made repeated measurements of edges, a necessary prerequisite for estimating edge error of the type considered here.

One case in which it is common to make repeated measurements is that of social networks reconstructed from passive observations. In these studies one attempts to deduce an underlying network structure from records of interaction or proximity between individuals, records that are typically an imperfect reflection of the true network. As an example, we consider the "reality mining" study of Eagle and Pentland [33], in which a group of university students carried mobile phones that used special software to record when they were in proximity with each other. We take as our data set the observations made each Wednesday in March and April of 2005. (We choose weekly observations to remove weekly periodic effects, and March and April because they fall during the university term.) This gives us eight sets of measurements of the entire network, one set for each day, in which an edge means that two individuals were in physical proximity at some time during that day. Thus N_{ij} takes the same value of 8 for all node pairs in this case, while E_{ij} spans the full range from 0 to 8. The measurements are sparse—most pairs of individuals never meet—and vary significantly from one week to the next.

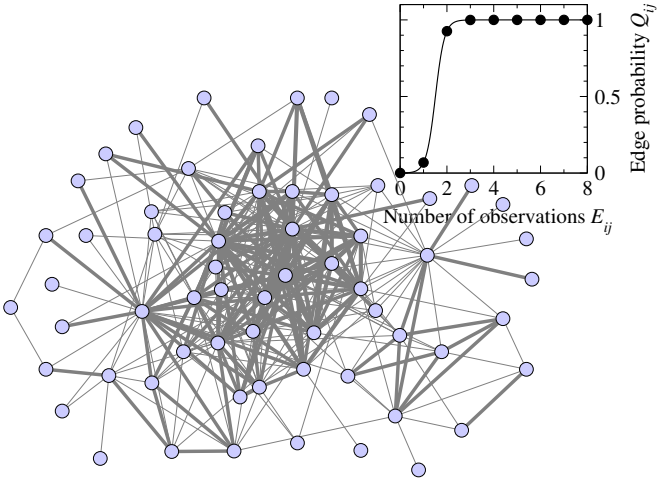

FIG. 1: Inferred ground-truth network for the reality mining data set. Edge widths indicate the inferred probabilities Q_{ij}. Edges with probability equal to the minimum possible value ($Q_{ij} = 0.0004$, corresponding to $E_{ij} = 0$ in Eq. (9)), are omitted, as are singleton nodes with no observed connections. Inset: Edge probability as a function of the number of observations E_{ij}.

Applying the methods of this paper to these data, the EM algorithm converges rapidly and reliably to parameter values $\alpha = 0.4242$, $\beta = 0.0043$, and $\rho = 0.0335$. These numbers tell us first that the network is indeed sparse (the value of ρ is small) and that there are very few false positives (the value of β implies that an edge is observed where none exists less than 1% of the time). On the other hand, even if the false positive rate is low, the probability of being wrong when one does observe an edge can still be high. This probability, which is given by $\beta(1-\rho)/[\alpha\rho+\beta(1-\rho)]$, is equal to 0.2270 in the present case, meaning that more than one in every five observed edges is in error. Moreover, the relatively small value of α implies that there are also a large number of false negatives: around 58% of pairs of individuals who are in fact connected in the underlying network are not observed in proximity on any one day. This is understandable. Most people do not see all of their acquaintances every day.

Figure 1 shows the inferred ground-truth network with edge widths indicating the probabilities Q_{ij}. The figure reveals a dense core of about twenty nodes that are with high probability connected to one another and a sparser periphery of nodes for whom the surety of connection is much lower. The thickest edges shown have $Q_{ij} > 0.999$, while the thinnest have $Q_{ij} < 0.1$. Inset at top right in the figure we also show the inferred relationship between the number of observations E_{ij} of a particular edge and the posterior probability Q_{ij}. As the figure shows, an edge observed only zero or one times implies a low Q_{ij}, so a single observation is probably a false alarm. But two or more observations of the same edge results in a value of

Q_{ij} approaching 1, indicating a strong inference that the edge exists in the ground truth. This makes it possible to infer the presence of edges with high probability despite the high false negative rate.

We have here worked though one example of an error model in detail, and we believe this particular model to be a good starting point for error estimation in many cases. It is simple and intuitive and, while not always wholly realistic, it can nonetheless usefully capture uncertainty in many structural measurements in the same way that the independent Gaussian model of conventional statistics, while often unrealistic, usefully captures error in real-valued data.

There are, however, many other possible error models one could use. The fundamental equations for the EM algorithm, Eqs. (3) and (4), can in principle be applied to any model to estimate errors. One can imagine modest modifications of the model used here to make it more realistic, such as allowing for different true and false positive rates in different parts of a network, perhaps depending on additional metadata or hidden variables describing node or edge types. However, one can also imagine completely different error models that could be useful in some cases [13], such as models for affiliation networks in which edges represent common membership of groups, or models for fixed choice social surveys in which respondents are limited in the number of connections they can report. In this short paper we have also neglected the issue of node error mentioned in the introduction. In principle, the methods described here could be extended to node errors as well, but we leave those developments for future work.

The author thanks George Cantwell, Travis Martin, Gesine Reinert, and Maria Riolo for useful comments. This work was funded in part by the US National Science Foundation under grant DMS–1407207.

[1] P. D. Killworth and H. R. Bernard, Informant accuracy in social network data. *Human Organization* **35**, 269–286 (1976).

[2] H. R. Bernard and P. D. Killworth, Informant accuracy in social network data II. *Human Communications Research* **4**, 3–18 (1977).

[3] P. V. Marsden, Network data and measurement. *Annual Review of Sociology* **16**, 435–463 (1990).

[4] C. von Mering, L. J. Jensen, B. Snel, S. D. Hooper, M. Krupp, M. Foglierini, N. Jouffre, M. A. Huynen, and P. Bork, STRING: Known and predicted protein-protein associations, integrated and transferred across organisms. *Nucleic Acids Research* **33**, D433–D437 (2005).

[5] N. J. Krogan *et al.*, Global landscape of protein complexes in the yeast Saccharomyces cerevisiae. *Nature* **440**, 637–643 (2006).

[6] A. Lakhina, J. Byers, M. Crovella, and P. Xie, Sampling biases in IP topology measurements. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies*, Institute of Electrical and Electronics Engineers, New York (2003).

[7] A. Clauset and C. Moore, Accuracy and scaling phenomena in Internet mapping. *Phys. Rev. Lett.* **94**, 018701 (2005).

[8] S. L. Feld and W. C. Carter, Detecting measurement bias in respondent reports of personal networks. *Social Networks* **24**, 365–383 (2002).

[9] S. P. Borgatti, K. M. Carley, and D. Krackhardt, On the robustness of centrality measures under conditions of imperfect data. *Social Networks* **28**, 124–136 (2006).

[10] G. Kossinets, Effects of missing data in social networks. *Social Networks* **28**, 247–268 (2006).

[11] D. J. Wang, X. Shi, D. A. McFarland, and J. Leskovec, Measurement error in network data: A reclassification. *Social Networks* **34**, 396–409 (2012).

[12] N. Erman and L. Todorovski, The effects of measurement error in case of scientific network analysis. *Scientometrics* **104**, 453–473 (2015).

[13] C. T. Butts, Network inference, error, and informant (in)accuracy: A Bayesian approach. *Social Networks* **25**, 103–140 (2003).

[14] M. S. Handcock and K. J. Gile, Modeling social networks from sampled data. *Annals of Applied Statistics* **4**, 5–25 (2010).

[15] D. Lusher, J. Koskinen, and G. Robins, *Exponential Random Graph Models for Social Networks: Theory, Methods, and Applications*. Cambridge University Press, Cambridge (2012).

[16] D. Liben-Nowell and J. Kleinberg, The link-prediction problem for social networks. *J. Assoc. Inf. Sci. Technol.* **58**, 1019–1031 (2007).

[17] A. Clauset, C. Moore, and M. E. J. Newman, Hierarchical structure and the prediction of missing links in networks. *Nature* **453**, 98–101 (2008).

[18] R. Guimerà and M. Sales-Pardo, Missing and spurious interactions and the reconstruction of complex networks. *Proc. Natl. Acad. Sci. USA* **106**, 22073–22078 (2009).

[19] M. Huisman, Imputation of missing network data: Some simple procedures. *Journal of Social Structure* **10**, 1–29 (2009).

[20] M. Kim and J. Leskovec, The network completion problem: Inferring missing nodes and edges in networks. In B. Liu, H. Liu, C. Clifton, T. Washio, and C. Kamath (eds.), *Proceedings of the 2011 SIAM International Conference on Data Mining*, pp. 47–58, Society for Industrial and Applied Mathematics, Philadephia, PA (2011).

[21] N. R. Smalheiser and V. I. Torvik, Author name disambiguation. *Annual Review of Information Science and Technology* **43**, 287–313 (2009).

[22] C. A. D'Angelo, C. Giuffrida, and G. Abramo, A heuristic approach to author name disambiguation in bibliometrics databases for large-scale research assessments. *J. Assoc. Inf. Sci. Technol.* **62**, 257–269 (2011).

[23] A. A. Ferreira, M. A. Goncalves, and A. H. F. Laender, A brief survey of automatic methods for author name disambiguation. *SIGMOD Record* **41**, 15–26 (2012).

[24] J. Tang, A. C. M. Fong, B. Wang, and J. Zhang, A unified probabilistic framework for name disambiguation in digital library. *IEEE Transactions on Knowledge and Data Engineering* **24**, 975–987 (2012).

[25] X. Han, Z. Shen, W.-X. Wang, and Z. Di, Robust reconstruction of complex networks from sparse data. *Phys.*

Rev. Lett. **114**, 028701 (2015).

[26] A. P. Dempster, N. M. Laird, and D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm. *J. R. Statist. Soc. B* **39**, 185–197 (1977).

[27] G. J. McLachlan and T. Krishnan, *The EM Algorithm and Extensions*. Wiley-Interscience, New York, 2nd edition (2008).

[28] B. Ball, B. Karrer, and M. E. J. Newman, An efficient and principled method for detecting communities in networks. *Phys. Rev. E* **84**, 036103 (2011).

[29] M. Evans, I. Guttman, Y. Haitovsky, and T. Swartz, Bayesian analysis of binary data subject to misclassification. In D. Berry, K. M. Chaloner, and J. Geweke (eds.), *Bayesian Statistics and Econometrics*, pp. 66–77, Wiley, New York (1995).

[30] W. H. Batchelder and A. K. Romney, Test theory without an answer key. *Psychometrika* **53**, 71–92 (1988).

[31] Z. Oravecz, R. Anders, and W. H. Batchelder, Hierarchical Bayesian modeling for test theory without an answer key. *Psychometrika* **86**, 341–364 (2015).

[32] T. Martin, B. Ball, and M. E. J. Newman, Structural inference for uncertain networks. *Phys. Rev. E* **93**, 012306 (2016).

[33] N. Eagle and A. Pentland, Reality mining: Sensing complex social systems. *Journal of Personal and Ubiquitous Computing* **10**, 255–268 (2006).